

Hyper-Spectral Image Analysis with Partially-Latent Regression and Spatial Markov Dependencies

Antoine Deleforge[†] Florence Forbes^{*} Sileye Ba^{*} Radu Horaud^{*}

^{*} INRIA Grenoble Rhône-Alpes, Montbonnot Saint-Martin, France

[†] Friedrich-Alexander-Universität, Erlangen-Nürnberg, 91058 Erlangen, Germany

Abstract—Hyper-spectral data can be analyzed to recover physical properties at large planetary scales. This involves resolving inverse problems which can be addressed within machine learning, with the advantage that, once a relationship between physical parameters and spectra has been established in a data-driven fashion, the learned relationship can be used to estimate physical parameters for new hyper-spectral observations. Within this framework, we propose a spatially-constrained and partially-latent regression method which maps high-dimensional inputs (hyper-spectral images) onto low-dimensional responses (physical parameters such as the local chemical composition of the soil). The proposed regression model comprises two key features. Firstly, it combines a Gaussian mixture of locally-linear mappings (GLLiM) with a partially-latent response model. While the former makes high-dimensional regression tractable, the latter enables to deal with physical parameters that cannot be observed or, more generally, with data contaminated by experimental artifacts that cannot be explained with noise models. Secondly, spatial constraints are introduced in the model through a Markov random field (MRF) prior which provides a spatial structure to the Gaussian-mixture hidden variables. Experiments conducted on a database composed of remotely sensed observations collected from the Mars planet by the Mars Express orbiter demonstrate the effectiveness of the proposed model.

Index Terms—Hyper-spectral images; Mars physical properties; OMEGA instrument; High-dimensional regression; Mixture models; Latent variable model, Markov random field.

I. INTRODUCTION

In modern geosciences, the use of remotely sensed data acquired from airborne, spacecraft, or satellite sensors allows to acquire multimodal observations at planetary levels [1]–[4]. Nowadays, hyper-spectral observations are used to monitor geophysical and environmental phenomena characterized by complex spectral signatures covering wide spectrum ranges.

Usually, recovering physical parameters such as the local chemical composition of the soil from hyper-spectral images involves solving inverse problems. This typically requires building complex high-dimensional to low-dimensional mappings relating the hyper-spectral space to the physical parameter spaces. Many approaches can be considered such as physically-based models, which build high-to-low dimensional mappings from geophysical models [5]–[9]. These models follow the laws of physical phenomena and are driven by cause-to-effect relationships. In [5], vegetation reflectance models

are built to extract vegetation variables from satellite sensors directional and spectral information. [6], [7] integrate hyper-spectral observations to a canopy radiative transfer model for forest species classification, and forest biophysical and biochemical properties estimation. [8], [9] use multi-spectral observations from multiple sensors, combined with physical radiative transfer models, to estimate ice surface characteristics such as temperature, concentration, and thickness. There are issues related to physically-based models, namely: expert knowledge is required for their construction. Moreover, due to strong non-linear interactions between a large number of variables, their inversion is not straightforward. Indeed, inverse problems are in general ill-posed in the presence of real data. Thus, every model requires a carefully designed inversion procedure.

Another family of approaches learn input-output mappings in a data driven fashion. These approaches have several advantages. Firstly, they do not require complete knowledge of the underlying physical phenomena. Therefore, they can be applied whenever training data and reliable models are available. Secondly, learned models can be readily applied to newly sensed data [10]. Furthermore, in situations where expert knowledge is available, learning based approaches can benefit from physics-based models, *e.g.*, by exploiting training data produced by radiative transfer models [11]. Machine learning methods are very popular for hyper-spectral data classification in the case of discrete-valued output variables. The reader is referred to [12] for an up to date review of state-of-the-art classification models applied to hyper-spectral data. Recently, [13] proposed a deep learning approach for hyper-spectral data classification. When the output variables are continuous, existing regression methods are mainly concerned with hyper-spectral data un-mixing where the goal is to identify end-members constituents together with their corresponding abundance from hyper-spectral observed mixtures [14], [15].

In the machine learning literature, regression models have been extensively studied. When dealing with high-dimensional input data, such as hyper-spectral observations, the vast majority of methods start by reducing the dimensionality followed by regression [16]–[18]. Such two-step approaches cannot be conveniently expressed in terms of a single optimization problem. Moreover, this presents the risk to map the input onto an intermediate low-dimensional space that does not necessarily contain the information needed to correctly predict

the output. To estimate non-linear mappings, mixtures of locally linear models were proposed [19]–[21], but not in the case of high-dimensional inputs and in the presence of partially-observed responses. An alternative popular approach consists in the use of kernel functions, with the drawbacks that these functions cannot be appropriately chosen automatically, and that the learned mappings cannot be inverted [17], [22]–[24].

In this paper we propose a learning based regression model that maps high-dimensional hyper-spectral observations onto low-dimensional geophysical parameters. The proposed model generalizes the Gaussian mixture of locally-linear mapping (GLLiM) model recently proposed [25]. The design principles of GLLiM consist of first, learning the parameters of an *inverse* low-to-high dimensional probabilistic mixture of locally-linear affine transformations, and second, deriving the *forward* posterior conditional density of the response, given an observed input and the learned model parameters. Exchanging the roles of the input and of the response variables during training, makes high-dimensional regression tractable and the conditional density of the response has a closed-form solution. Moreover, the model is designed such that it can deal with a partially-latent response: the vector-valued response variable is composed of both observed entries and latent (unobserved) entries. This enables to handle in a principled fashion the presence, in the training data, of physical parameters and experimental artifacts that cannot be explicitly measured.

With respect to [25], the main novelty of the model proposed in this paper is the introduction of a Markov random field (MRF) prior on the hidden (assignment) variables of the mixture model. This prior explicitly enforces spatial regularity of the reconstructed geophysical fields. Our model integrates high-dimensional to low-dimensional regression and MRF spatial regularization into a unified Bayesian framework. Because model parameter learning is intractable due to the complex Markov interactions between the variables, in particular because of the MRF prior, the parameters are estimated via an approximate variational Bayesian learning scheme. We provide a detailed description of a variational EM (VEM) procedure that is used for learning the model parameters from training data. Our modeling is quite original as in general MRF priors are enforced onto response variable fields [14], [26]. Furthermore, because the mixture's hidden variables are discrete and finite, while response variables are continuous, our setting is computationally more tractable. Another interesting feature is that the Markov field is directly included in the generative modeling of the high-to-low dimensional mapping. In general, such dimension reduction aims at representing spatial vectors in a lower more tractable space in which noise can be better characterized. Indeed, the dimension of the space spanned by spectra from an image is generally much lower than the available number of bands. When the goal is to characterize observed spectra with a limited number of features (such as abundances or physical characteristics), identifying appropriate subspaces can then significantly improve performance. This reduction step is then often an intermediate step before additional refinement such as the inclusion

of contextual information is performed. Typically, in [27], dimension reduction and Markov modeling are carried out in two successive steps and assuming a linear mixing model. This is quite different from our approach which aims at modeling and identifying non-linear mappings and not necessarily in the un-mixing context only. In our work, the high-to-low mapping is a target while it is only a convenient intermediate step in most studies. Of course, considering the high level of activity in hyper-spectral image analysis, it would be challenging to present an exhaustive view of the subject but we believe that methods using subspace reduction in this field correspond to a different use of dimensionality reduction. We refer to [28] for one recent review on both subspace identification and spatial modeling.

The model that we propose is also related to hyper-spectral data un-mixing [14], [28], [29] since in our model the constituting components of the hyper-spectral data are segregated with a locally-affine regression model. However and in contrast to existing methods, our model, through the presence of the partially-latent output, is able to handle the presence of non identifiable components in the input data. This particular property could be interpreted as an ability to partially un-mix hyper-spectral observations while letting a remaining part unexplained.

The effectiveness of the proposed model is assessed using datasets of simulated and real hyper-spectral images of the OMEGA imaging spectrometer on board of the Mars Express spacecraft [1]. Thorough quantitative and qualitative results demonstrate that our model outperforms several state-of-the-art regression models. Furthermore, the fact that we integrate an MRF prior produces spatially-regular parameter maps.

The remainder of this paper is organized as follows. Section II summarizes the high-dimensional partially-latent regression model [25]. Section III presents the novel MRF spatially-constrained regression model. Section IV applies the proposed method to hyper-spectral images collected from Mars with the OMEGA instrument. Finally, Section V draws conclusions and discusses possible future research challenges.

II. INVERSE REGRESSION WITH PARTIALLY-LATENT RESPONSE VARIABLES

In this section, we summarize the inverse regression model proposed in [25]. This model maps low-dimensional variables $\mathbf{X} = \{\mathbf{X}_n\}_{n=1}^N$ onto high-dimensional variables $\mathbf{Y} = \{\mathbf{Y}_n\}_{n=1}^N$. To each low dimensional variable $\mathbf{X}_n \in \mathbb{R}^L$ corresponds a high dimensional variable $\mathbf{Y}_n \in \mathbb{R}^D$ with $L \ll D$. We consider the *hybrid*¹ *Gaussian locally-linear mapping* (hGLLiM) framework proposed in [25] that treats each \mathbf{X}_n as a partially latent variable, namely:

$$\mathbf{X}_n = \begin{bmatrix} \mathbf{T}_n \\ \mathbf{W}_n \end{bmatrix}, \quad (1)$$

¹hybrid refers to the observed and latent structure of the response variable \mathbf{X}_n .

where $\mathbf{T}_n \in \mathbb{R}^{L_t}$ is observed and $\mathbf{W}_n \in \mathbb{R}^{L_w}$ is latent, with $L = L_t + L_w$. Let us consider $\{\mathbf{x}_n\}_{n=1}^N$ and $\{\mathbf{y}_n\}_{n=1}^N$, where $(\mathbf{y}_n, \mathbf{x}_n)$ is a realization of $(\mathbf{Y}_n, \mathbf{X}_n)$, such that \mathbf{y}_n is the noisy image of \mathbf{x}_n obtained from a K -component locally-affine transformation. This is modeled by introducing the latent variable $Z_n \in \{1, \dots, K\}$ such that:

$$\mathbf{Y}_n = \sum_{k=1}^K \mathbb{I}(Z_n = k) (\mathbf{A}_k \mathbf{X}_n + \mathbf{b}_k + \mathbf{E}_k) \quad (2)$$

where \mathbb{I} is the indicator function, matrix $\mathbf{A}_k \in \mathbb{R}^{D \times L}$ and vector $\mathbf{b}_k \in \mathbb{R}^D$ define an affine transformation and $\mathbf{E}_k \in \mathbb{R}^D$ is an error term capturing both the observation noise in \mathbb{R}^D and the reconstruction error due to the local affine approximation. Under the assumption that \mathbf{E}_k is a zero-mean Gaussian variable with covariance matrix $\Sigma_k \in \mathbb{R}^{D \times D}$ that does not depend on \mathbf{X}_n , \mathbf{Y}_n , and Z_n , we obtain:

$$p(\mathbf{Y}_n = \mathbf{y}_n | \mathbf{X}_n = \mathbf{x}_n, Z_n = k; \theta) = \mathcal{N}(\mathbf{y}_n; \mathbf{A}_k \mathbf{x}_n + \mathbf{b}_k, \Sigma_k), \quad (3)$$

where θ is the set of all the model parameters. It has to be noticed that the roles of the input and of the output (or response) variables are interchanged, *i.e.*, the low-dimensional variable \mathbf{X}_n is the regressor.

To complete the definition and enforce the affine transformations to be local, each \mathbf{X}_n is assumed to follow a mixture of K Gaussians defined by

$$p(\mathbf{X}_n = \mathbf{x}_n | Z_n = k; \theta) = \mathcal{N}(\mathbf{x}_n; \mathbf{c}_k, \Gamma_k) \quad (4)$$

with means $\mathbf{c}_k \in \mathbb{R}^L$ and covariance matrices $\Gamma_k \in \mathbb{R}^{L \times L}$. In the standard hGLLiM model, a multinomial prior

$$p(Z_n = k; \theta) = \pi_k \quad (5)$$

is used on the hidden variable Z_n . An efficient closed-form EM algorithm [25], provides maximum likelihood estimates of the parameters $\theta = \{\pi_k, \mathbf{c}_k, \Gamma_k, \mathbf{A}_k, \mathbf{b}_k, \Sigma_k\}$ given the observed training data $\{\mathbf{y}_n, \mathbf{t}_n\}_{n=1}^N$.

A. Inverse Regression Strategy

Once the parameter vector θ has been estimated, one obtains an inverse regression from \mathbb{R}^L (low-dimensional space) to \mathbb{R}^D (high-dimensional space), using the following *inverse conditional density*:

$$p(\mathbf{Y}_n = \mathbf{y}_n | \mathbf{X}_n = \mathbf{x}_n; \theta) = \sum_{k=1}^K \frac{p(Z_n = k) \mathcal{N}(\mathbf{x}_n; \mathbf{c}_k, \Gamma_k)}{\sum_{j=1}^K p(Z_n = j) \mathcal{N}(\mathbf{x}_n; \mathbf{c}_j, \Gamma_j)} \mathcal{N}(\mathbf{y}_n; \mathbf{A}_k \mathbf{x}_n + \mathbf{b}_k, \Sigma_k). \quad (6)$$

The forward regression of interest, *i.e.*, from \mathbb{R}^D (the high dimension) to \mathbb{R}^L (the low dimension), is obtained from the *forward conditional density*:

$$p(\mathbf{X}_n = \mathbf{x}_n | \mathbf{Y}_n = \mathbf{y}_n; \theta) = \sum_{k=1}^K \frac{p(Z_n = k) \mathcal{N}(\mathbf{y}_n; \mathbf{c}_k^*, \Gamma_k^*)}{\sum_{j=1}^K p(Z_n = j) \mathcal{N}(\mathbf{y}_n; \mathbf{c}_j^*, \Gamma_j^*)} \mathcal{N}(\mathbf{x}_n; \mathbf{A}_k^* \mathbf{y}_n + \mathbf{b}_k^*, \Sigma_k^*). \quad (7)$$

The latter involves the *forward regression parameters*:

$$\theta^* = \{\mathbf{c}_k^*, \Gamma_k^*, \mathbf{A}_k^*, \mathbf{b}_k^*, \Sigma_k^*\}_{k=1}^K \quad (8)$$

that can be analytically derived from the *inverse regression parameters* θ that were previously learned:

$$\pi_k^* = \pi_k, \quad (9)$$

$$\mathbf{c}_k^* = \mathbf{A}_k \mathbf{c}_k + \mathbf{b}_k, \quad (10)$$

$$\Gamma_k^* = \Sigma_k + \mathbf{A}_k \Gamma_k \mathbf{A}_k^\top, \quad (11)$$

$$\mathbf{A}_k^* = \Sigma_k^* \mathbf{A}_k^\top \Sigma_k^{-1}, \quad (12)$$

$$\mathbf{b}_k^* = \Sigma_k^* (\Gamma_k^{-1} \mathbf{c}_k - \mathbf{A}_k^\top \Sigma_k^{-1} \mathbf{b}_k), \quad (13)$$

$$\Sigma_k^* = (\Gamma_k^{-1} + \mathbf{A}_k^\top \Sigma_k^{-1} \mathbf{A}_k)^{-1}. \quad (14)$$

Learning an inverse regression model to obtain the desired forward mapping is referred to as *inverse regression strategy*. This is motivated by a drastic reduction of the model's size, making tractable its estimation. Indeed, if we consider that the error vectors \mathbf{E}_k are modeled by isotropic Gaussian noise with equal variances, the dimension of the learned parameter vector θ is $\mathcal{O}(DL + L^2)$, while it would be $\mathcal{O}(DL + D^2)$ using a forward model².

B. Estimating the Latent Dimension with BIC

The proposed model involves latent response variables \mathbf{W}_n of dimensionality L_w . The *Bayesian information criterion* (BIC) can be used to select a value for the latent dimension. Once a set of model parameters $\tilde{\theta}$ has been learned, BIC can be computed as follows:

$$BIC(\tilde{\theta}, N) = -2\mathcal{L}(\tilde{\theta}) + \mathcal{D}(\tilde{\theta}) \log N, \quad (15)$$

where \mathcal{L} denotes the observed-data log-likelihood and $\mathcal{D}(\tilde{\theta})$ denotes the dimension of the complete parameter vector $\tilde{\theta}$. Assuming isotropic and equal noise covariance matrices $\{\Sigma_k\}_{k=1}^K$, we have:

$$\mathcal{D}(\tilde{\theta}) = K(D(L_w + L_t + 1) + L_t(L_t + 3)/2 + 1), \quad (16)$$

$$\mathcal{L}(\tilde{\theta}) = \sum_{n=1}^N \log p(\mathbf{y}_n, \mathbf{t}_n; \tilde{\theta}), \quad (17)$$

A natural way to define a value for L_w is, for a given value of K , to train the model with different values of L_w , and select the value minimizing BIC. We will refer to the corresponding method as hGLLiM-BIC. While being computationally demanding, the BIC parameter selection procedure has the advantage of not requiring manually setting L_w .

III. VARIATIONAL EM FOR THE SPATIAL MARKOV MODEL

In this section we propose and derive an extension of hGLLiM that accounts for spatially located variables, *i.e.*, \mathbf{Y}_n is defined at the sites of hyper-spectral data, where each site $n \in S = \{1 \dots N\}$ typically corresponds to a graph vertex or to a pixel. The novelty of the model proposed in this paper with respect to [25] lies in the explicit spatial dependencies which

²Recall that $L \ll D$.

are modeled by introducing the assumption that variables \mathbf{Z} are distributed according to a discrete Markov random field (MRF) on S with pairwise potentials. The multinomial prior (5) is therefore replaced by:

$$p(\mathbf{Z} = \mathbf{z}; \boldsymbol{\psi}) = \mathcal{K}(\boldsymbol{\psi})^{-1} \exp(H(\mathbf{z}; \boldsymbol{\psi})) \quad (18)$$

with H the energy function $H(\mathbf{z}; \boldsymbol{\psi}) = \sum_{n=1}^N \left(\alpha_{z_n} + \frac{\beta}{2} \sum_{m \in \Omega(n)} \mathbb{I}(z_n = z_m) \right)$ and the normalizing constant $\mathcal{K}(\boldsymbol{\psi}) = \sum_{\mathbf{z}} \exp(H(\mathbf{z}; \boldsymbol{\psi}))$. In these expressions, $\Omega(n)$ denotes the neighbors of site n in S (typically the vertices linked to n by an edge), $\boldsymbol{\psi} = \{\boldsymbol{\alpha}, \beta\}$, $\boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_K)$ is a vector of external field parameters and β is a scalar interaction parameter. In practice, a constraint on $\boldsymbol{\alpha}$ is required to avoid identifiability issues. Thus, we assume $\alpha_1 = 0$. We also assume that variables from different site locations $\{(\mathbf{X}_n, \mathbf{Y}_n)\}_{n=1}^N$ are conditionally independent given $\{Z_n\}_{n=1}^N$. Therefore, the parameters of this model are $\phi = \{\boldsymbol{\theta}, \boldsymbol{\psi}\}$ defined by:

$$\boldsymbol{\theta} = \{\mathbf{c}_k, \boldsymbol{\Gamma}_k, \mathbf{A}_k, \mathbf{b}_k, \boldsymbol{\Sigma}_k\}_{k=1}^K \quad (19)$$

$$\boldsymbol{\psi} = \{\boldsymbol{\alpha}, \beta\}. \quad (20)$$

Because \mathbf{X}_n defined in (1) is a *partially-latent* variable, the parameter estimation procedure exploits observed pairs $\{\mathbf{y}_n, \mathbf{t}_n\}_{n=1}^N$ while being constrained by the presence of the latent variables \mathbf{W}_n . The decomposition of \mathbf{X}_n into observed and latent components implies that the model parameters \mathbf{c}_k , $\boldsymbol{\Gamma}_k$ and \mathbf{A}_k must be decomposed as well. Assuming the independence of \mathbf{T} and \mathbf{W} given \mathbf{Z} , this gives:

$$\mathbf{c}_k = \begin{bmatrix} \mathbf{c}_k^t \\ \mathbf{c}_k^w \end{bmatrix}, \boldsymbol{\Gamma}_k = \begin{bmatrix} \boldsymbol{\Gamma}_k^t & \mathbf{0} \\ \mathbf{0} & \boldsymbol{\Gamma}_k^w \end{bmatrix}, \mathbf{A}_k = \begin{bmatrix} \mathbf{A}_k^t & \mathbf{A}_k^w \end{bmatrix}.$$

Considering the complete data, with observed variables $(\mathbf{Y}_n, \mathbf{T}_n)_{1:N}$ and hidden variables $(Z_n, \mathbf{W}_n)_{1:N}$, the corresponding expectation-maximisation (EM) algorithm is intractable due to the complex Markovian dependencies between variables Z_n . Therefore, we resort to a variational approximation which consists of approximating the true posterior $p(\mathbf{w}, \mathbf{z} | \mathbf{y}, \mathbf{t}, \phi)$ by a distribution $q_{\mathbf{W}, \mathbf{Z}}(\mathbf{w}, \mathbf{z})$ which factorizes as

$$q_{\mathbf{W}, \mathbf{Z}}(\mathbf{w}, \mathbf{z}) = \prod_{n=1}^N q_{\mathbf{W}_n, Z_n}(\mathbf{w}_n, z_n).$$

From this factorization follows a variational EM (VEM) parameter learning procedure which consists of N successive E-steps denoted by $E(\mathbf{W}_n, Z_n)$ for $n = \{1, \dots, N\}$, followed by variational M-steps.

A. Variational E- (\mathbf{W}_n, Z_n) step

Given current parameters values $\phi^{(i)}$ and distribution $q_{\mathbf{W}, \mathbf{Z}}^{(i)}$, the E- (\mathbf{W}_n, Z_n) step consists of computing the probability distribution $q_{\mathbf{W}_n, Z_n}^{(i+1)}(\mathbf{w}_n, z_n)$ which is proportional to:

$$\exp \left(E_{q_{\mathbf{W} \setminus n, \mathbf{Z} \setminus n}^{(i)}} [\log p(\mathbf{w}_n, z_n | \mathbf{y}, \mathbf{t}, \mathbf{W} \setminus n, \mathbf{Z} \setminus n; \phi^{(i)})] \right) \quad (21)$$

where $\mathbf{W} \setminus n$ (resp. $\mathbf{Z} \setminus n$) denotes the variables $\{\mathbf{W}_m, m = 1 : N, m \neq n\}$ (resp. $\{Z_m, m = 1 : N, m \neq n\}$).

It follows from (21) that we can also write:

$$q_{\mathbf{W}_n, Z_n}^{(i+1)}(\mathbf{w}_n, z_n) = q_{\mathbf{W}_n | Z_n}^{(i+1)}(\mathbf{w}_n | z_n) q_{Z_n}^{(i+1)}(z_n)$$

where:

$$q_{Z_n}^{(i+1)}(z_n = k) = \frac{p(\mathbf{y}_n, \mathbf{t}_n | z_n = k; \boldsymbol{\theta}^{(i)}) \exp(\alpha_k^{(i)} + \beta^{(i)} \sum_{m \in \Omega(n)} q_{Z_m}^{(i)}(k))}{\sum_{l=1}^K p(\mathbf{y}_n, \mathbf{t}_n | z_n = l; \boldsymbol{\theta}^{(i)}) \exp(\alpha_l^{(i)} + \beta^{(i)} \sum_{m \in \Omega(n)} q_{Z_m}^{(i)}(l))} \quad (22)$$

$$q_{\mathbf{W}_n | Z_n}^{(i+1)}(\mathbf{w}_n | z_n = k) = \mathcal{N}(\mathbf{w}_n; \tilde{\boldsymbol{\mu}}_{nk}^w, \tilde{\boldsymbol{\Sigma}}_k^w) \quad (23)$$

with:

$$\tilde{\boldsymbol{\mu}}_{nk}^w = \tilde{\boldsymbol{\Sigma}}_k^w ((\mathbf{A}_k^w)^{\top} (\boldsymbol{\Sigma}_k^{(i)})^{-1} (\mathbf{y}_n - \mathbf{A}_k^t \mathbf{t}_n - \mathbf{b}_k^{(i)}) + (\boldsymbol{\Gamma}_k^{(i)})^{-1} \mathbf{c}_k^w) \quad (24)$$

$$\tilde{\boldsymbol{\Sigma}}_k^w = ((\boldsymbol{\Gamma}_k^{(i)})^{-1} + (\mathbf{A}_k^w)^{\top} (\boldsymbol{\Sigma}_k^{(i)})^{-1} \mathbf{A}_k^w)^{-1}. \quad (25)$$

B. Variational M-Steps

The variational M-steps consist in estimating the parameters $\phi = \{\boldsymbol{\theta}, \boldsymbol{\psi}\}$ that maximize the expectation:

$$E_{q_{\mathbf{W}, \mathbf{Z}}^{(i+1)}} [\log p(\mathbf{y}, \mathbf{t}, \mathbf{W}, \mathbf{Z}; \boldsymbol{\theta}, \boldsymbol{\psi})]. \quad (26)$$

It follows that the update of parameters in $\boldsymbol{\theta}$ is identical to that of the standard hGLLiM EM [25], where the posterior densities $p(Z_n = k | \mathbf{y}_n, \mathbf{t}_n; \boldsymbol{\theta}^{(i)})$ are now replaced by the variational distributions $q_{Z_n}^{(i+1)}(z_n = k)$. As in [25], the update of noise covariance matrices $\{\boldsymbol{\Sigma}_k\}_{k=1}^K$ depends on the way they are constrained. In practice, we enforce isotropic covariances for all k in order to avoid over-parameterization.

Once $\boldsymbol{\theta}$ has been updated, the proposed MRF model induces an additional M-step for the update of $\boldsymbol{\psi} = \{\boldsymbol{\alpha}, \beta\}$. This step does not admit an explicit closed-form expression but can be solved numerically using gradient descent schemes. It is straightforward to show that the maximization of (26) in $\boldsymbol{\psi}$ admits a unique solution. Indeed, it is equivalent to solve

$$\begin{aligned} \boldsymbol{\psi}^{(i+1)} &= \arg \max_{\boldsymbol{\psi}} E_{q_{\mathbf{Z}}^{(i+1)}} [\log p(\mathbf{Z}; \boldsymbol{\psi})] \\ &= \arg \max_{\boldsymbol{\psi}} E_{q_{\mathbf{Z}}^{(i+1)}} [H(\mathbf{z}; \boldsymbol{\psi})] - \log \mathcal{K}(\boldsymbol{\psi}). \end{aligned} \quad (27)$$

Denoting the gradient vector and Hessian matrix respectively by $\nabla_{\boldsymbol{\psi}}$ and $\nabla_{\boldsymbol{\psi}}^2$, it comes

$$\begin{aligned} \nabla_{\boldsymbol{\psi}} E_{q_{\mathbf{Z}}^{(i+1)}} [\log p(\mathbf{Z}; \boldsymbol{\psi})] &= E_{q_{\mathbf{Z}}^{(i+1)}} [\nabla_{\boldsymbol{\psi}} H(\mathbf{Z}; \boldsymbol{\psi})] \\ &= E_{p(\mathbf{z}; \boldsymbol{\psi})} [\nabla_{\boldsymbol{\psi}} H(\mathbf{Z}; \boldsymbol{\psi})] \end{aligned} \quad (28)$$

$$\begin{aligned} \nabla_{\boldsymbol{\psi}}^2 E_{q_{\mathbf{Z}}^{(i+1)}} [\log p(\mathbf{Z}; \boldsymbol{\psi})] &= E_{q_{\mathbf{Z}}^{(i+1)}} [\nabla_{\boldsymbol{\psi}}^2 H(\mathbf{Z}; \boldsymbol{\psi})] \\ &= E_{p(\mathbf{z}; \boldsymbol{\psi})} [\nabla_{\boldsymbol{\psi}}^2 H(\mathbf{Z}; \boldsymbol{\psi})] \\ &= \text{var}_{p(\mathbf{z}; \boldsymbol{\psi})} [\nabla_{\boldsymbol{\psi}} H(\mathbf{Z}; \boldsymbol{\psi})] \end{aligned} \quad (29)$$

The last expectations in (28) and (29) are taken over the Potts prior (18). It follows that whenever $H(\mathbf{z}; \boldsymbol{\psi})$ is linear in $\boldsymbol{\psi}$,

$\nabla_{\psi}^2 H(\mathbf{Z}; \psi)$ is zero, the Hessian matrix is a semi-definite negative matrix and the function to optimize is concave. Unfortunately, due to the intractable normalizing constant \mathcal{K} , expressions (28) and (29) are not directly available. It is necessary to approximate the terms involving the true MRF prior $p(\mathbf{z}; \psi)$ using an approximation. A natural approach is to use:

$$q_{\mathbf{Z}}^{\text{prior}}(\mathbf{z}; \alpha, \beta) = \prod_{n=1}^N q_{Z_n}^{\text{prior}}(z_n; \alpha, \beta) \quad (30)$$

with $q_{Z_n}^{\text{prior}}(z_n; \alpha, \beta)$ defined by:

$$q_{Z_n}^{\text{prior}}(z_n = k; \alpha, \beta) = \frac{\exp(\alpha_k + \beta \sum_{m \in \Omega(n)} q_{Z_m}^{(i)}(k))}{\sum_{l=1}^K \exp(\alpha_l + \beta \sum_{m \in \Omega(n)} q_{Z_m}^{(i)}(l))} \quad (31)$$

This MRF prior approximation induced by the posterior variational approximation has been proposed in [30] and also exploited in [31]. It follows that the gradient in (28) can be approximated by (with $\alpha_1 = 0$):

$$\begin{cases} \frac{1}{2} \sum_{n=1}^N \sum_{m \in \Omega(n)} \sum_{k=1}^K (q_{Z_n}^{(i+1)}(k) q_{Z_m}^{(i+1)}(k) \\ - q_{Z_n}^{\text{prior}}(k; \alpha, \beta) q_{Z_m}^{\text{prior}}(k; \alpha, \beta)) \\ \sum_{n=1}^N (q_{Z_n}^{(i+1)}(2) - q_{Z_n}^{\text{prior}}(2; \alpha, \beta)) \\ \vdots \\ \sum_{n=1}^N (q_{Z_n}^{(i+1)}(K) - q_{Z_n}^{\text{prior}}(K; \alpha, \beta)) \end{cases} \quad (32)$$

C. Inverse Regression Strategy with Markov Random Fields

With this proposed MRF extension of hGLLiM, the derivation of the desired forward conditional density (7) is less straightforward. Indeed, while equations (7)-(8) remain valid, they now involve the marginal of the MRF prior $p(Z_n = k)$ which is known to be hard to compute (see (18)). Approximations are required and a natural candidate is $q_{Z_n}^{\text{prior}}(k)$ which has to be computed iteratively using the variational EM algorithm with θ (or θ^*) held fixed. Note that α could also either be set fixed to the learned values or re-estimated. At iteration (i), for each n , given current values for $\alpha^{(i)}, \beta^{(i)}$ and $q_{Z_m}^{(i)}(k), q_{Z_n}^{(i)}(k)$ is updated using:

$$q_{Z_n}^{(i+1)}(z_n = k) = \frac{p(\mathbf{y}_n | z_n = k; \theta^*) \exp(\alpha_k^{(i)} + \beta^{(i)} \sum_{m \in \Omega(n)} q_{Z_m}^{(i)}(k))}{\sum_{l=1}^K p(\mathbf{y}_n | z_n = l; \theta^*) \exp(\alpha_l^{(i)} + \beta^{(i)} \sum_{m \in \Omega(n)} q_{Z_m}^{(i)}(l))} \quad (33)$$

Although the update formula for $q_{Z_n}(k)$ is synchronous, an asynchronous update is known to be more efficient. Then parameters $\alpha^{(i+1)}, \beta^{(i+1)}$ are obtained by solving (27) as explained in Section III-B. At algorithm convergence, $p(Z_n = k)$

is set as:

$$p(Z_n = k) \approx q_{Z_n}^{\text{prior}}(k; \alpha^{(\infty)}, \beta^{(\infty)}) = \frac{\exp(\alpha_k^{(\infty)} + \beta^{(\infty)} \sum_{m \in \Omega(n)} q_{Z_m}^{(\infty)}(k))}{\sum_{l=1}^K \exp(\alpha_l^{(\infty)} + \beta^{(\infty)} \sum_{m \in \Omega(n)} q_{Z_m}^{(\infty)}(l))} \quad (34)$$

When required, a response \mathbf{x}_n corresponding to an input \mathbf{y}_n can be obtained using the expectation of $p(\mathbf{X}_n = \mathbf{x}_n | \mathbf{Y}_n = \mathbf{y}_n)$ that can be approximated using:

$$\mathbb{E}[\mathbf{X}_n | \mathbf{y}_n; \theta^*, \psi] \approx \sum_{k=1}^K \frac{q_{Z_n}^{\text{prior}}(k; \alpha^{(\infty)}, \beta^{(\infty)}) \mathcal{N}(\mathbf{y}_n; \mathbf{c}_k^*, \mathbf{\Gamma}_k^*)}{\sum_{j=1}^K q_{Z_n}^{\text{prior}}(j; \alpha^{(\infty)}, \beta^{(\infty)}) \mathcal{N}(\mathbf{y}_n; \mathbf{c}_j^*, \mathbf{\Gamma}_j^*)} (\mathbf{A}_k^* \mathbf{y}_n + \mathbf{b}_k^*). \quad (35)$$

D. Prediction with spatial constraints: MRF-hGLLiM

To complete this section, we provide an overview of the proposed procedure implementing the inverse regression strategy while accounting for spatial interactions. It divides into two steps, a learning step followed by a prediction step. Each step is summarized below:

Spatial inverse-regression learning:

- *Training observed data:* $\{(\mathbf{y}_n, \mathbf{t}_n)\}_{n=1}^N$;
- *Missing variables:* $\{(\mathbf{W}_n, \mathbf{Z}_n)\}_{n=1}^N$;
- *Model:* MRF-hGLLiM (hybrid Gaussian locally linear mapping with spatial constraints) defined by (1), (3), 4, and (18));
- *Parameters:* $\theta = \{\mathbf{c}_k, \mathbf{\Gamma}_k, \mathbf{A}_k, \mathbf{b}_k, \mathbf{\Sigma}_k\}_{k=1}^K$ and $\psi = \{\alpha, \beta\}$;
- *Estimation procedure:* VEM (section III);
- *Output:*
 - $\hat{\theta}$: estimated inverse regression parameters θ ,
 - $\hat{\theta}^*$: estimated forward regression parameters θ^* , defined by (10) to (14),
 - $\hat{\psi}$: estimated Markov parameters ψ .

Spatial forward-prediction:

- *Observed data:* $\{\tilde{\mathbf{y}}_n\}_{n=1}^{\tilde{N}}$ (no observed $\{\tilde{\mathbf{t}}_n\}_{n=1}^{\tilde{N}}$);
- *Missing variables:* $\{(\mathbf{X}_n, \mathbf{Z}_n)\}_{n=1}^{\tilde{N}}$ where $\mathbf{X}_n = [\mathbf{T}_n; \mathbf{W}_n]$;
- *Model:* MRF-hGLLiM with θ fixed to $\hat{\theta}$;
- *Estimation procedure:* VEM (section III) with no M- θ step (θ fixed to $\hat{\theta}$) which reduces to the procedure described in section III-C;
- *Output:*
 - $\tilde{\psi} = \{\tilde{\alpha}, \tilde{\beta}\}$,
 - Estimated predictions for $\{\mathbf{x}_n\}_{n=1}^{\tilde{N}}$, i.e., $\tilde{\mathbf{x}}_n = \mathbb{E}[\mathbf{X}_n | \tilde{\mathbf{y}}_n; \hat{\theta}^*, \tilde{\psi}]$ computed using approximation (35).

IV. RETRIEVING MARS PHYSICAL PROPERTIES FROM HYPER-SPECTRAL IMAGES

In order to evaluate the proposed model, we use a database composed of synthetic spectra with their associated physical

parameter values. This database was generated using the radiative transfer model presented in [11] in order to investigate hyperspectral images collected from the imaging spectrometer OMEGA [1] onboard of the Mars Express spacecraft. The synthetic dataset is composed of 15,407 spectra associated with five physical parameter values, namely:

- (i) *proportion of water ice* (Prop. H₂O),
- (ii) *proportion of CO₂ ice* (Prop. CO₂),
- (iii) *proportion of dust* (Prop. Dust),
- (iv) *grain size of water ice* (Size H₂O), and
- (v) *grain size of CO₂ ice* (Size CO₂).

Each spectrum is made of 184 wavelengths. Compared to other available hyper-spectral image databases, *e.g.* [32], [33], this one combines several desirable features, making it suitable for the evaluation of the proposed model for the following reasons. First, the radiative transfer model can be used to generate as much data as required to reliably estimate the model parameters. Second, generated spectra present a highly non-linear relationship with underlying physical parameters. Third, the associated real hyper-spectral images of Mars surface contain spatial dependencies between neighboring pixels.

A. Hybrid-GLLiM Evaluation on Synthetic Individual Spectra

In order to numerically evaluate the performances of MRF-hGLLiM on the synthetic data, we adopted a cross-validation strategy: one part of the synthetic database is used for training and the remaining part is used for testing. Note that because all these spectra were generated independently, there are no spatial dependencies present in the training dataset. Therefore, the value of the MRF interaction parameter β was set to 0 in that case, *i.e.*, no neighborhood dependencies are accounted for during training. With this setting, the proposed MRF-hGLLiM and hGLLiM models are equivalent. The synthetic data are used, firstly to learn an inverse *low-dimensional to high-dimensional regression function* between physical parameters and spectra from the database, and secondly to estimate the unknown physical parameters corresponding to a newly observed individual spectrum and using the forward mapping (35), where $q_{Z_n}^{\text{prior}}(k)$ is replaced by e^{α_k} , or equivalently π_k .

The hGLLiM algorithm was compared to JGMM (joint Gaussian mixture model) [20], MLE (mixture of linear experts) [19], SIR (sliced inverse regression) [16] and RVM (multivariate relevance vector machine) [24]. When a value $L_w > 0$ is used for training, we denote by hGLLiM- L_w the hybrid GLLiM algorithm. As shown in [25], JGMM and MLE are equivalent to hGLLiM-0 with covariance matrices $\{\Sigma_k\}_{k=1}^K$ being respectively unconstrained and diagonal. SIR is used with one (SIR-1) or two (SIR-2) principal axes for dimensionality reduction, 20 slices³, and polynomial regression of order three⁴. SIR quantizes the low-dimensional data \mathbf{X} into *slices* or clusters which in turn induces a quantization of the \mathbf{Y} -space. Each \mathbf{Y} -slice (all points \mathbf{y}_n that map to the same

\mathbf{X} -slice) is then replaced with its mean and PCA is carried out on these means. The resulting dimensionality reduction is then informed by \mathbf{X} values through the preliminary slicing. RVM may be viewed as a multivariate probabilistic formulation of *support vector regression* [34]. As with any kernel method, RVM critically depends on the choice of a kernel function. Using the authors' publicly available software⁵, we ran preliminary tests to determine an optimal kernel choice for each dataset under consideration. We tested 14 kernels with 10 different scales, ranging from 1 to 30, resulting in a total of 140 kernels for each dataset.

As already mentioned an objective evaluation is performed using cross-validation. A total of 10,000 training input-output couples are randomly selected twenty times as the training set, and the 5,407 remaining spectra are used for testing. For all models, the training data are normalized to zero mean and unit variance. Normalization is then reversed on test data and estimated output to obtain the final results. This normalization showed to significantly improve the performances of all methods. We use $K = 50$ for MLE, hGLLiM and JGMM. MLE and JGMM are constrained with isotropic covariance matrices as this parametrization achieves the best results. For RVM, a third degree polynomial kernel with scale 6 is selected as the top performing among 140 tested kernels using cross-validation.

As a quality measure of the estimated parameters, we compute the normalized root mean-square error (NRMSE) which quantifies the difference between estimated and ground-truth parameter values \hat{t}_m and t_m :

$$\text{NRMSE} = \sqrt{\frac{\sum_{m=1}^M (\hat{t}_m - t_m)^2}{\sum_{m=1}^M (t_m - \bar{t})^2}} \text{ with } \bar{t} = \frac{1}{M} \sum_{m=1}^M t_m. \quad (36)$$

NRMSE is normalized thus enabling direct comparison between the parameters which are of very different range: closer to zero, more accurate the predicted values.

a) *Fully Observed Output*: Using these different regression models, from 184-dimensional spectra, we start by retrieving the five parameter values, namely, proportion of water ice (Prop. H₂O), proportion of CO₂ ice (Prop. CO₂), proportion of dust (Prop. Dust), grain size of water ice (Size H₂O), and grain size of CO₂ ice (Size CO₂). The training is done with synthetic spectra annotated with the five parameters, and hence output variables are fully observed. Results obtained with the five methods are summarized in Table I. MLE and hGLLiM perform similarly and outperform the three other methods in estimating each parameter. The average training time of MLE was 12.6 seconds and the average training time of hGLLiM was 18.9 seconds using our Matlab implementation. In this task, as expected, using 1 or 2 additional latent components in hGLLiM only shows a slight improvement compared to MLE, since the output is fully observed during training. Notice that the obtained mean NRMSE with the proportion of water (Table I, column 2) and the grain size of CO₂ (Table I, column 6) parameters are very large, *i.e.*, above 0.5 for all methods. This

³The number of slices is known to have very little influence on the results.

⁴Higher orders did not show significant improvements in our experiments.

⁵http://www.mvrv.com/Multivariate_Relevance_Vector

TABLE I

NORMALIZED ROOT MEAN SQUARED ERROR (NRMSE) FOR MARS SURFACE PHYSICAL PROPERTIES RECOVERED FROM INDIVIDUAL HYPER-SPECTRAL VECTORS, USING SYNTHETIC DATA, DIFFERENT METHODS AND FULLY-OBSERVED-OUTPUT TRAINING.

Method	Prop. H ₂ O	Prop. CO ₂	Prop. Dust	Size H ₂ O	Size CO ₂
JGMM	2.40 ± 18.5	0.84 ± 1.64	0.63 ± 1.02	0.73 ± 1.02	1.08 ± 4.52
SIR-1	3.41 ± 20.0	1.28 ± 2.16	1.04 ± 1.79	0.69 ± 0.92	1.85 ± 7.24
SIR-2	3.27 ± 18.6	0.96 ± 1.75	0.89 ± 1.53	0.62 ± 0.86	1.66 ± 6.53
RVM	1.28 ± 7.57	0.50 ± 0.95	0.40 ± 0.69	0.51 ± 0.67	0.89 ± 3.80
MLE (hGLLiM-0)	1.04 ± 6.66	0.37 ± 0.72	0.28 ± 0.50	0.45 ± 0.74	0.60 ± 2.59
hGLLiM-1	0.95 ± 5.92	0.34 ± 0.65	0.24 ± 0.44	0.42 ± 0.71	0.56 ± 2.44
hGLLiM-2	0.99 ± 6.02	0.36 ± 0.70	0.27 ± 0.48	0.40 ± 0.66	0.58 ± 2.66

TABLE II

NORMALIZED ROOT MEAN SQUARED ERROR (NRMSE) FOR MARS SURFACE PHYSICAL PROPERTIES RECOVERED FROM INDIVIDUAL HYPER-SPECTRAL VECTORS, USING SYNTHETIC DATA, DIFFERENT METHODS AND PARTIALLY-OBSERVED-OUTPUT TRAINING.

Method	Proportion of CO ₂ ice	Proportion of dust	Grain size of H ₂ O ice
JGMM	0.83 ± 1.61	0.62 ± 1.00	0.79 ± 1.09
SIR-1	1.27 ± 2.09	1.03 ± 1.71	0.70 ± 0.94
SIR-2	0.96 ± 1.72	0.87 ± 1.45	0.63 ± 0.88
RVM	0.52 ± 0.99	0.40 ± 0.64	0.48 ± 0.64
MLE (hGLLiM-0)	0.54 ± 1.00	0.42 ± 0.70	0.61 ± 0.92
hGLLiM-1	0.36 ± 0.70	0.28 ± 0.49	0.45 ± 0.75
hGLLiM-2*†	0.34 ± 0.63	0.25 ± 0.44	0.39 ± 0.71
hGLLiM-3	0.35 ± 0.66	0.25 ± 0.44	0.39 ± 0.66
hGLLiM-4	0.38 ± 0.71	0.28 ± 0.49	0.38 ± 0.65
hGLLiM-5	0.43 ± 0.81	0.32 ± 0.56	0.41 ± 0.67
hGLLiM-20	0.51 ± 0.94	0.38 ± 0.65	0.47 ± 0.71
hGLLiM-BIC	0.34 ± 0.63	0.25 ± 0.44	0.39 ± 0.71




Fig. 1. Ground truth and estimated dust concentration maps with MRF-hGLLiM and hGLLiM using synthetic hyper-spectral images. For visualization purpose, the image color dynamics have been adjusted to emphasize the differences.

suggests that the relationship between these parameters and observed spectra is complex and hard to learn.

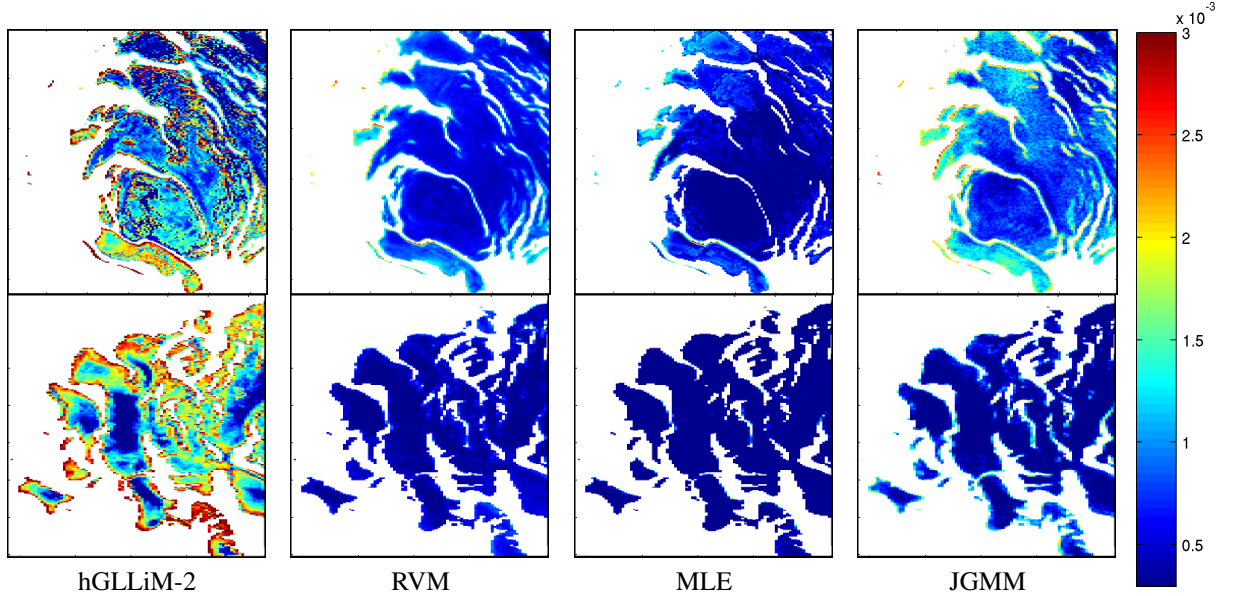
b) Partially Latent Output: In order to fully illustrate the potential of hybrid GLLiM, we deliberately ignore the values of two of the five parameters in the database and treat these two parameters as *latent variables*. We chose to ignore the proportion of water ice and the grain size of CO₂ ice since they yield the worst reconstruction errors when they are fully observed (see Table I). A previous study [10] reveals that these two parameters are sensitive to the same wavelengths and are suspected to *mix* with the other three parameters in the radiative transfer model used to generate the synthetic data. Therefore, not only that they are difficult to estimate, but they also seem to affect the estimation of the three other parameters.

Table II shows NRMSE values obtained for the three remaining parameters. The ground-truth latent variable dimension is $L_w^* = 2$, and accordingly, the empirically best latent dimension for hGLLiM is $L_w^\dagger = 2$. hGLLiM-2 outperforms all the other methods on that task, more precisely the error is 36% lower than the second best-performing method, RVM, closely followed by MLE. No significant difference is observed between hGLLiM-2 and hGLLiM-3. The average training time of MLE was 12.2 seconds and the average training time of hGLLiM was 18.8 seconds using our Matlab implementation. Due to the computation of the $D \times D$ kernel matrix, the training time of RVM was about 10 times larger.


We also tested the selection of the latent dimension (L_w)

TABLE III
NORMALIZED ROOT MEAN SQUARED ERROR (NRMSE) OF hGLLiM-0 AND MRF-GLLiM FOR MARS SURFACE PHYSICAL PROPERTIES RECOVERED FROM HYPER-SPECTRAL IMAGES, USING SYNTHETIC DATA AND FULLY-OBSERVED-OUTPUT TRAINING.

Method	Prop. H ₂ O	Prop. CO ₂	Prop. Dust	Size H ₂ O	Size CO ₂
hGLLiM-0	0.27 ± 0.02	0.25 ± 0.03	0.25 ± 0.03	0.29 ± 0.05	0.24 ± 0.04
MRF-hGLLiM-0	0.25 ± 0.03	0.23 ± 0.03	0.23 ± 0.03	0.26 ± 0.05	0.23 ± 0.04



(a) Proportion of dust



(b) Proportion of CO₂ ice

Fig. 2. Proportions obtained with our method (hGLLiM-2) and three other methods. The data correspond to hyper-spectral images grabbed from two different viewpoints of the South polar cap of Mars. Top rows: orbit 41; Bottom rows: orbit 61. White areas correspond to unexamined regions, where the synthetic model does not apply.




Fig. 3. Qualitative results obtained with MRF-hGLLiM-2 for orbit 41. Top row: Each pixel n is colored according to the class k maximizing $p(Z_n = k | \mathbf{Y}_n = \mathbf{y}_n; \phi)$ (each class index is assigned an arbitrary color). Middle row: Proportion of CO₂ ice. Bottom row: Proportion of Dust. Concentration color codes are the same as in Fig. 2. To better appreciate the smoothing effect, a zoomed region is shown.




Fig. 4. Qualitative results obtained with MRF-hGLLiM-2 for orbit 61. Top row: Each pixel n is colored according to the class k maximizing $p(Z_n = k | \mathbf{Y}_n = \mathbf{y}_n; \phi)$ (each class index is assigned an arbitrary color). Middle row: Proportion of CO₂ ice. Bottom row: Proportion of Dust. Concentration color codes are the same as in Fig. 2. To better appreciate the smoothing effect, a zoomed region is showed for each orbit.

based on BIC (Section II-B). For each training set, the hGLLiM-BIC method minimizes BIC with the latent dimension in the range $0 \leq L_w \leq 20$, and used the corresponding model to perform the regression. Interestingly, hGLLiM-BIC performed very well on these large training sets ($N = 10,000$) as it correctly selects $L_w = 2$ for the 20 training sets used by BIC (the BIC selection could differ with the training set), yielding the same results as those obtained with hGLLiM-2. Interestingly, notice how hGLLiM and almost all the methods benefit from removing the “faulty” parameters that are difficult to estimate. In all cases, a slight performance increase for the reconstruction of the remaining three others parameters is observed when comparing to the results of Table I.

B. MRF-hGLLiM Evaluation on Synthetic Hyper-spectral Images

The data used in the previous section is not spatial, in the sense that there is no dependency between the different vectors of observations. In order to assess the advantage of the proposed MRF extension of hGLLiM when spatial dependencies exist, we synthesized 50 hyper-spectral images from the synthetic spectra dataset presented in the previous section. We generated 400×300 pixel images consisting of 12 (4×3) square regions (see Figure 1), where each region corresponds to a set of similar underlying physical parameters. More specifically, for each region, a vector of

physical parameters in \mathbb{R}^{L_t} was picked, and the region was filled with this vector as well as its 15 nearest neighbors in the dataset. A synthetic hyper-spectral image was then generated from associated spectra, and corrupted with additive white Gaussian noise with a signal-to-noise-ratio of 6 dB.

We first learned the model parameters $\hat{\theta}$ using hGLLiM-0 (equivalent to MLE [19]) and 10,000 individual spectra associated to all 5 physical parameters, as explained in the previous section. We then compared the performance of MRF-hGLLiM-0 when β was automatically estimated from the test images (section III-C) and when β was set to 0, *i.e.* no spatial dependency, which corresponds to hGLLiM-0. The spectra used in the 50 test synthetic images were outside the training set. The MRF prior was incorporated using a 8-pixel neighborhood and MRF-hGLLiM-0 estimated β values comprised between 2.8 et 3.1 for these synthetic images. Table III shows the NRMSE obtained with both methods. The use of spatial dependencies significantly reduced the estimation errors for all concentrations, in all 50 test images. For each image, the averaged across pixels errors were computed for both hGLLiM-0 and MRF-hGLLiM-0 and the corresponding two samples of size 50 were compared using a paired t-test. The test confirmed a highly significant gain in favor of MRF-hGLLiM-0. The average computational time for a 100×100 test image is 220s using hGLLiM-0 and 550 using MRF-hGLLiM-0 with our Matlab implementation.

C. Real Mars Hyper-spectral Images

We now investigate the potential of the proposed MRF-hGLLiM method using a dataset of real hyper-spectral images collected from the imaging spectrometer OMEGA instrument [1] on board of the Mars Express spacecraft. We adopt the evaluation protocol proposed in [10]. An adequately selected subset of 6,983 spectra of the synthetic database is used for model training. Then, trained models are applied to the real data made of observed spectra acquired from satellite hyper-spectral sensors. For evaluation, we focus on data from Mars's South polar cap. Since ground-truth for the physical properties of Mars's South pole regions is not currently available, we consider a qualitative evaluation by comparing five models: hGLLiM-2, MRF-hGLLiM-2, and the three best performing methods, among the tested ones, namely RVM, MLE and JGMM.

First, hGLLiM-2, RVM, MLE and JGMM are used to extract physical parameters from two hyper-spectral images of the same area but seen from two slightly different view points (orbit 41 and orbit 61). Since responses are proportions with values between 0 and 1, values smaller than 0 or higher than 1 are not admissible and hence they are set to either one of the closest bounds. In the real data setting, the best performance is achieved when no data mean-variance normalization is performed, and using $K = 10$ for hGLLiM-2, MLE and JGMM. As it can be seen in Fig. 2, hGLLiM-2 estimates proportion maps with similar characteristics for the two view points, which suggests a good model-response

consistency. Such a consistency is not observed with the other three tested methods. Furthermore, RVM and MLE result in a much larger number of values outside the admissible interval $[0, 1]$. Moreover, hGLLiM-2 is the only method featuring less dust at the South pole cap center and higher proportions of dust at the boundaries of the CO2 ice, which matches expected results from planetology [35]. Finally, note that the proportions of CO2 ice and dust clearly seem to be complementary using hGLLiM-2, while this complementarity is not obviously noticed when using the other methods.

Despite these satisfying qualitative results, hGLLiM-2 produces less regular maps than the other methods. This is particularly true for orbit 41, as can be seen in the top rows of Fig. 2(a) and (b). This does not match physical expectations since chemical proportions should vary rather smoothly in space. To address this issue, we use MRF-hGLLiM-2 by exploiting the MRF-forward mapping (35).

The MRF-hGLLiM-2 model is trained using synthetic data without accounting for the spatial dependencies (β set to 0). This training results in a model equivalent to hGLLiM-2. However, during testing, the MRF prior is enforced using an 8-pixel neighborhood. For illustration purpose, the parameter β is empirically set to control spatial smoothness of the resulting physical parameter map. According to the MRF prior model, the larger the value of β , the smoother the resulting parameter field.

We compare the outputs of hGLLiM-2 and MRF-hGLLiM-2 with different interaction parameter values β . The top rows of Fig. 3 and Fig. 4 color each pixel according to the class k maximizing $p(Z_n = k | \mathbf{Y}_n = \mathbf{y}_n; \phi)$. By progressively increasing β from 0 to 500, isolated points are removed, thus creating larger and larger image regions associated to a single class. As it can be seen in Fig. 3, this has the desired smoothing effect on the resulting proportion maps. Interestingly, MRF-hGLLiM-2 yields smoother maps than hGLLiM-2 while preserving its desirable qualitative performances, *e.g.*, higher proportions of dust at boundaries and complementarity of CO2 and dust proportions. Notice that β has a less significant effect on proportions obtained at orbit 61, *e.g.*, Fig. 4, since hGLLiM-2 already yields quite smooth maps for this orbit.

V. CONCLUSIONS

In this paper, we proposed a Gaussian mixture of locally-affine regression model that maps a high-dimensional space of hyper-spectral observations onto a low-dimensional space of physical parameters. Spatial regularity is enforced and integrated into the model through an MRF prior on the Gaussian-mixture hidden variables. Model learning is achieved via a variational expectation-maximization procedure that is fully derived and described in detail. Model evaluation is conducted using both simulated data and real hyper-spectral images from the Mars Express satellite sensors. The proposed approach outperforms four state-of-the-art regression methods. Furthermore, when applied to actual images of the Mars South pole, our model produces spatially-regular and smooth maps in accordance with the MRF constraints.

As a first model extension, we will investigate its application to the hyper-spectral un-mixing problem. In this context, we will design experimental setups to study the ability of the model to partially estimate hyper-spectral mixture components together with their corresponding proportions. Further research, in line with dynamical geophysical data-assimilation research [36], [37], will be to investigate generalizations of the model in order to process temporal sequences of observations. To this end, relevant dynamical models of physical phenomena driving hyper-spectral observation evolution will be investigated and integrated. Equipped with an MRF spatial prior and a temporal dynamical model, our formulation could be able to address missing-data interpolation problems which are ubiquitous in remote sensing [38].

REFERENCES

- [1] J.-P. Bibring *et al.*, "OMEGA: Observatoire pour la Minéralogie, l'Eau, les Glaces et l'Activité," in *Mars Express: The Scientific Payload*, vol. 1240, 2004, pp. 37–49.
- [2] G. Chang *et al.*, "The new age of hyperspectral oceanography," *Oceanography*, vol. 17, no. 2, pp. 16–23, 2004.
- [3] E. Cloutis, B. Vila, A. Bell, and M. Lamothe, "Hyperspectral and luminescence observer (HALO) Mars mission concept - innovative data triage, compression, processing and analysis for the hyperspectral imager," *IEEE Workshop on Hyperspectral Image and Signal Processing: Evolution in Remote Sensing*, vol. 93, no. 3, pp. 1–4, 2009.
- [4] F. Hilton *et al.*, "Hyperspectral earth observation from IASI: Five years of accomplishments," *Bulletin of the American Meteorological Society*, vol. 93, no. 3, pp. 347–370, 2012.
- [5] D. S. Kimes, Y. Knyazikhin, J. L. Privette, A. A. Abuelgasim, and F. Gao, "Inversion methods for physically based models," *Remote Sensing Reviews*, vol. 18, no. 2–4, pp. 381–439, 2000.
- [6] D. Goodenough, B. Victoria, J. Li, G. Asner, M. Schaepman, L. Ustin, and A. Dyk, "Combining hyperspectral remote sensing and physical modeling for applications in land ecosystems," in *IEEE International Conference on Geoscience and Remote Sensing*, 2006, pp. 2000–2004.
- [7] P. Stenberg, M. Mottus, and M. Rautiainen, "Modeling the spectral signature of forests: Application of remote sensing models to coniferous canopies," *Springer Advances in Land Remote Sensing*, pp. 147–171, 2008.
- [8] M. Shokr, "A physics-based remote sensing data fusion approach," in *International Conference on Geoscience and Remote Sensing Symposium*. IEEE, 2003, pp. 1050–1052.
- [9] L. Foster, B. Brock, M. Cutler, and F. Diotri, "A physically based method for estimating supraglacial debris thickness from thermal band remote-sensing data," *Journal of Glaciology*, vol. 58, no. 210, pp. 677–691, 2012.
- [10] C. Bernard-Michel, S. Douté, M. Fauvel, L. Gardes, and S. Girard, "Retrieval of Mars surface physical properties from OMEGA hyperspectral images using regularized sliced inverse regression," *Journal of Geophysical Research: Planets*, vol. 114, no. E6, 2009.
- [11] S. Douté, E. Deforas, F. Schmidt, R. Oliva, and B. Schmitt, "A comprehensive numerical package for the modeling of Mars hyperspectral images," in *Lunar and Planetary Science XXXVIII*, March 2007.
- [12] M. Fauvel, Y. Tarabalka, J. A. Benediktsson, J. Chanussot, and J. C. Tilton, "Advances in spectral-spatial classification of hyperspectral images," *Proceedings of the IEEE*, vol. 101, no. 3, pp. 652–675, 2013.
- [13] Y. Chen, Z. Lin, X. Zhao, G. Wang, and Y. Gu, "Deep learning-based classification of hyperspectral data," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 7, no. 3, pp. 2094–2107, 2014.
- [14] J. Bhatt, M. Joshi, and M. Raval, "A data-driven stochastic approach for unmixing hyperspectral imagery," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 7, no. 6, pp. 1936–1946, 2014.
- [15] R. Heylen, M. Parente, and P. Gader, "A review of nonlinear hyperspectral unmixing methods," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 7, no. 6, pp. 1844–1868, 2014.
- [16] K. C. Li, "Sliced inverse regression for dimension reduction," *Journal of the American Statistical Association*, vol. 86, no. 414, pp. 316–327, 1991.
- [17] H. Wu, "Kernel sliced inverse regression with applications to classification," *Journal of Computational and Graphical Statistics*, vol. 17, no. 3, pp. 590–610, 2008.
- [18] K. P. Adraghi and R. D. Cook, "Sufficient dimension reduction and prediction in regression," *Philosophical Transactions of the Royal Society A*, vol. 367, no. 1906, pp. 4385–4405, 2009.
- [19] L. Xu, M. I. Jordan, and G. E. Hinton, "An alternative model for mixtures of experts," in *Advances in Neural Information Processing Systems*. MIT Press, Cambridge, MA, USA, December 1995, pp. 633–640.
- [20] Y. Qiao and N. Minematsu, "Mixture of probabilistic linear regressions: A unified view of GMM-based mapping techniques," in *IEEE International Conference on Acoustics, Speech, and Signal Processing*, April 2009, pp. 3913–3916.
- [21] S. Ingrassia, S. C. Minotti, and G. Vittadini, "Local statistical modeling via a cluster-weighted approach with elliptical distributions," *Journal of classification*, vol. 29, no. 3, pp. 363–401, 2012.
- [22] M. Tipping, "Sparse Bayesian learning and the relevance vector machine," *The Journal of Machine Learning Research*, vol. 1, pp. 211–244, 2001.
- [23] N. Lawrence, "Probabilistic non-linear principal component analysis with gaussian process latent variable models," *The Journal of Machine Learning Research*, vol. 6, pp. 1783–1816, 2005.
- [24] A. Thayananthan, R. Navaratnam, B. Stenger, P. Torr, and R. Cipolla, "Multivariate relevance vector machines for tracking," in *European Conference on Computer Vision*. Springer, 2006, pp. 124–138.
- [25] A. Deleforge, F. Forbes, and R. Horaud, "High-Dimensional Regression with Gaussian Mixtures and Partially-Latent Response Variables," *Statistics and Computing*, 2014.
- [26] Y. Tarabalka, M. Fauvel, J. Chanussot, and J. Benediktsson, "SVM- and MRF-based method for accurate classification of hyperspectral images," *IEEE Geoscience and Remote Sensing Letters*, vol. 7, no. 4, pp. 736–740, 2010.
- [27] J. Li, J. M. Bioucas-Dias, and A. Plaza, "Spectral-Spatial Hyperspectral Image Segmentation Using Subspace Multinomial Logistic Regression and Markov Random Fields," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 50, no. 3, pp. 809–823, 2012.
- [28] J. M. Bioucas-Dias, A. Plaza, N. Dobigeon, M. Parente, Q. Du, P. Gader, and J. Chanussot, "Hyperspectral unmixing overview: Geometrical, statistical, and sparse regression-based approaches," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 5, no. 2, pp. 1–4, 2012.
- [29] Y. Altmann, N. Dobigeon, S. McLaughlin, and J.-Y. Tourneret, "Residual component analysis of hyperspectral images: Application to joint nonlinear unmixing and nonlinearity detection," *IEEE Transactions on Image Processing*, vol. 23, no. 5, pp. 2148–2158, 2014.
- [30] G. Celeux, F. Forbes, and N. Peyrard, "EM procedures using mean field-like approximations for Markov model-based image segmentation," *Pattern Recognition*, vol. 36, no. 1, pp. 131–144, 2003.
- [31] L. Chaari, T. Vincent, F. Forbes, M. Dojat, and P. Ciuciu, "Fast joint detection-estimation of evoked brain activity in event-related fMRI using a variational approach," *IEEE Transactions on Medical Imaging*, vol. 32, no. 5, pp. 821–837, May 2013.
- [32] J. Leenaars, "Africa soil profiles database, version 1.1. a compilation of geo-referenced and standardized legacy soil profile data for Sub-Saharan Africa (with dataset)," Africa Soil Information Service (AfSIS) project, Wageningen, the Netherlands, Tech. Rep. ISRIC report 2013/03, 2013.
- [33] L. Tits, B. Somers, J. Stuckens, and P. Coppin, "Validating nonlinear mixing models: benchmark datasets from vegetated areas," in *6th Workshop on Image and Signal Processing: Evolution in Remote Sensing (WHISPERS)*, June 2014.
- [34] A. J. Smola and B. Schölkopf, "A tutorial on support vector regression," *Statistics and computing*, vol. 14, no. 3, pp. 199–222, 2004.
- [35] S. Douté *et al.*, "Nature and composition of the icy terrains of the south pole of Mars from MEX OMEGA observations," in *Lunar and Planetary Science XXXVI*, March 2005.
- [36] Q. Sixian, M. Jianwen, and W. Xuanji, "Construction and experiment of hierarchical Bayesian network in data assimilation," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 6, no. 2, pp. 1036–1047, 2013.
- [37] Y. Dong, J. Wang, C. Li, G. Yang, Q. Wang, F. Liu, J. Zhao, H. Wang, and W. Huang, "Comparison and analysis of data assimilation algorithms for predicting the leaf area index of crop canopies," *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, vol. 6, no. 1, pp. 188–201, 2013.

- [38] A.-B. Salberg, "Land cover classification of cloud-contaminated multitemporal high-resolution images," *IEEE Transactions on Geoscience and Remote Sensing*, vol. 49, no. 1, pp. 377–387, 2011.



Antoine Deleforge received the B.Sc. (2008) and M.Sc. (2010) engineering degrees in computer science and mathematics from the Ecole Nationale Supérieure d'Informatique et Mathématiques Appliquées de Grenoble (ENSIMAG), France, as well as the specialized M.Sc. (2010) research degree in computer graphics, vision, robotics from the Université Joseph Fourier, Grenoble. In 2013, he received the Ph.D. degree in computer science and applied mathematics from the University of Grenoble. Since January 2014, he has as a post-doctoral

fellow appointment with the chair of Multimedia Communication and Signal Processing of the Erlangen-Nuremberg University, Germany. His research interests include machine learning for signal processing, Bayesian statistics, computational auditory scene analysis, and robot audition.



Florence Forbes received the B.Sc. and M.Sc. degrees in computer science and applied mathematics from the Ecole Nationale Supérieure d'Informatique et Mathématiques Appliquées de Grenoble (ENSIMAG), France, and the PhD degree in applied probabilities from the University Joseph Fourier, Grenoble, France. Since 1998, she has been a research scientist with the Institut National de Recherche en Informatique et Automatique (INRIA), Grenoble Rhône-Alpes, Montbonnot, France, where she founded the MISTIS team and has been the team

head since 2003. Her research activities include Bayesian analysis, Markov and graphical models, and hidden structure models.



Silève Ba received the M.Sc. (2000) in applied mathematics and signal processing from University of Dakar, Senegal, and the M.Sc. (2002) in mathematics, computer vision, and machine learning from Ecole Normale Supérieure de Cachan, Paris, France. From 2003 to 2009 he was a PhD student and then a post-doctoral researcher at IDIAP Research Institute, Martigny, Switzerland where he worked on probabilistic models for object tracking and human activity recognition. From 2009 to 2013, he was a researcher at Telecom Bretagne, Brest, France working on variational models for multi-modal geophysical data processing. From 2013 to 2014 he worked at RN3D Innovation Lab, Marseille, France, as a research engineer, where he used computer vision and machine learning principles and methods to develop human-computer interaction software tools. Since June 2014 he has been a researcher in the Perception team at INRIA Grenoble Rhône-Alpes working on machine learning and computer vision models for natural human-robot interaction.



Radu Horaud received the B.Sc. (1977) degree in electrical engineering, the M.Sc. (1979) degree in control engineering, and the Ph.D. (1981) degree in computer science from the Institut National Polytechnique de Grenoble, France. Currently he holds a position of director of research with the Institut National de Recherche en Informatique et Automatique (INRIA), Grenoble Rhône-Alpes, Montbonnot, France, where he is the founder and head of the PERCEPTION team. His research interests include computer vision, machine learning, audio signal processing, audiovisual analysis, and robotics and human-robot interaction. He is

an area editor of the *Computer Vision and Image Understanding*, a member of the advisory board of the *International Journal of Robotics Research*, and an associate editor of the *International Journal of Computer Vision*. In 2013, Radu Horaud was awarded a five year ERC Advanced Grant for his project *Vision and Hearing in Action* (VHIA, 2014-2019).